arXiv:1611.00864v1 [cs.NE] 3 Nov 2016

Recurrent Neural Networks for Spatiotemporal Dynamics of Intrinsic Networks from fMRI Data

R Devon Hjelm Sergey Plis Vince Calhoun

Functional magnetic resonance imaging (fMRI) of temporally-coherent blood oxygenization level-dependent (BOLD) signal provides an effective means of analyzing functionally coherent patterns in the brain [6, 5, 13]. Intrinsic networks [INs, 3] and functional connectivity are important outcomes of fMRI studies and are central to understanding brain function and making diagnoses [4, 1, 10]. The most popular method for separating INs, independent component analysis [ICA, 2], begins with the assumption that the data is a mixture of maximally independent sources. ICA is trainable through one of many relatively simple optimization routines that maximize non-Gaussianity or minimize mutual information [11]. Although fMRI data is a time series, ICA, as with other popular linear methods for separating INs, is order-agnostic in time: each multivariate signal at each time step is treated as i.i.d.. ICA in its common use in the field employs the same parameterization across subjects, which allows for either temporal or spatial variability, but not both. In order to overcome shortcomings of temporal ICA in lack of dynamics and subject-wise/temporal variability of spatial maps, but without abandoning the fundamental strengths of ICA, we combine recurrent neural networks (RNNs) with an ICA objective. The resulting model naturally represents temporal and spatial dynamics—having subject-wise and temporally variable spatial maps—and is easily trainable using gradient descent and back-propagation.

An RNN can represent an fMRI time series, $\mathbf{x}_{1:T,n} = (\mathbf{x}_{1,n}, \mathbf{x}_{2,n}, \dots \mathbf{x}_{T,n})$, through a product of conditional densities:

$$p(\mathbf{x}_{1:T,n}) = p(\mathbf{x}_{1,n}) \prod_{t=2}^{T} p(\mathbf{x}_{t,n} | \mathbf{x}_{1:t-1,n}), \tag{1}$$

where n indexes the full fMRI time series for a single subject. The RNN parameterizes the conditional dependence of $\mathbf{x}_{t,n}$ on the past, $\mathbf{x}_{1:t-1,n}$, via a set of deterministic hidden units, $\mathbf{h}_{1:t,n}$, which allow for simple feed-forward computation and gradient back-propagation. Define, as with ICA, a linear transformation for each observation to source configuration: $\mathbf{s}_{t,n} = \mathbf{W}\mathbf{x}_{t,n}$, and define a high-kurtosis and factorized source distribution, $p_{s_{t,n}}(\mathbf{s}_{t,n})$ (such as logistic or Laplace) for each time step, t, and each fMRI sequence, n. Let us apply the transformation to an fMRI time series: $\mathbf{s}_{1:T,n} = f(\mathbf{x}_{1:T,n}) = (\mathbf{W}\mathbf{x}_{1,n}, \mathbf{W}\mathbf{x}_{2,n}, \dots \mathbf{W}\mathbf{x}_{T,n})$. The log-likelihood function can be re-parameterized as:

$$\begin{aligned}
\log p(\mathbf{x}_{1:T,n}) &= \log p(\mathbf{x}_{1,n}) + \sum_{t=2}^{T} \log p(\mathbf{x}_{t,n} | \mathbf{x}_{1:t-1,n}) \\
&= \log p_{s_1}(\mathbf{W}\mathbf{x}_1) + \sum_{t=2}^{T} \log p_{s_{t,n}}(\mathbf{W}\mathbf{x}_{t,n} | \mathbf{x}_{1:t-1,n}) + T \log |\det \mathbf{J}_f(\mathbf{x}_{1:T,n})| \\
&= T \log |\det \mathbf{W}| + \log p_{s_{1,n}}(\mathbf{W}\mathbf{x}_{1,n}) + \sum_{t=2}^{T} \log p_{s_{t,n}}(\mathbf{W}\mathbf{x}_{t,n} | \mathbf{x}_{1:t-1,n}),
\end{aligned} \tag{2}$$

where $\mathbf{J}_f$ is the Jacobian and the source distribution, $p_{s_{t,n}}$, has parameters determined by $\mathbf{h}_{t,n}$. A high-kurtosis distribution is desirable to ensure independence of the sources (or minimizing the mutual information, e.g., the INFOMAX objective [2]), so a reasonable choice for the outputs of the RNN at each time step are the mean, $\boldsymbol{\mu}$, and scale, $\boldsymbol{\sigma}$, for a logistic distribution:

$$\boldsymbol{\mu}_{t,n} = \mathbf{W}_{\mu}\mathbf{h}_{t,n}; \qquad \boldsymbol{\sigma}_{t,n} = \mathbf{W}_{\sigma}\mathbf{h}_{t,n}. \tag{3}$$

Our treatment assumes the ICA weight matrix, $\mathbf{W}$, is square, which is necessary to ensure a tractable determinant Jacobian and inverse. fMRI data is very high dimensional, so to reduce the dimensionality of the necessarily square matrix we must resort to some sort of dimensionality reduction as preprocessing. Widely used in ICA studies of fMRI is principle component analysis (PCA), which can be used to reduce the dimensionality of the data, selecting number of components to match the dimension of the sources, $\mathbf{s}_{t,n}$.

The objective above is tractable in part because the transformation of the observations, f, is directional in dependence of $\mathbf{x}_t$ in time, which ensures that the Jacobian is upper diagonal in time. While the model is already novel, we can further leverage this directed dependence in the transformation to represent variable spatial maps across time and subjects.

Consider the transformation $\mathbf{s}_{t,n} = \mathbf{W}_{t,n}\mathbf{x}_{t,n}$, such that $\mathbf{W}_{t,n}$ is a subject-specific weight matrix at time t with form:

$$\mathbf{W}_{t,n} = \overline{\mathbf{W}} + \widehat{\mathbf{W}}_{t,n} = \overline{\mathbf{W}} + \mathbf{c}_{t,n}\mathbf{p}_{t,n}^T, \tag{4}$$

where $\widehat{\mathbf{W}}_{t,n} = \mathbf{c}_{t,n}\mathbf{p}_{t,n}^T$ is a rank-1 weight matrix composed of two equal-length vectors and $\overline{\mathbf{W}}$ is the same across time and subjects. To assure that the weights reflecting fMRI dynamics are subject-specific and are trainable using back-propagation, the factors of the rank-1 matrix, $\mathbf{c}$ and $\mathbf{p}$, are modeled as outputs of two separate neural networks with input $\mathbf{h}_t$. As the variable permutation to the spatial maps is rank-1, both the determinant of the Jacobian and inverse of a rank-1 matrix added to a full rank square matrix are easy to compute [12]:

$$\det(\mathbf{W}_{t,n}) = (1 + \mathbf{p}_{t,n}^T \overline{\mathbf{W}}^{-1}\mathbf{c}_{t,n})\det(\overline{\mathbf{W}}); \qquad \mathbf{W}_{t,n}^{-1} = \overline{\mathbf{W}}^{-1} - \frac{\overline{\mathbf{W}}^{-1}\mathbf{c}_{t,n}\mathbf{p}_{t,n}^T\overline{\mathbf{W}}^{-1}}{1 + \mathbf{p}_{t,n}^T\overline{\mathbf{W}}^{-1}\mathbf{c}_{t,n}}. \tag{5}$$

To demonstrate the properties and strengths of our model, we apply our method to fMRI task data. Data used in this work is comprised of task-related scans from 28 healthy participants and 24 subjects diagnosed with schizophrenia. All participants were scanned during an auditory oddball task (AOD) involving the detection of an infrequent target sound within a series of standard and novel sounds. More detailed information regarding participant demographics and task details is provided by Swanson et al. [14].

We used a simple RNN with a 100 recurrent hidden units and a recurrent parameterization and two feed forward networks with a single hidden layer to determine $\mathbf{c}_t$ and $\mathbf{p}_t$. The model was optimized via the RMSProp algorithm [8] and incorporated various weight decay strategies for regularization.

Our results indicate the model is able to separate interpretable and unique INs, with approximately 50 grey matter regions after filtering out white matter, ventricles, and well-known motion artifacts. Using the β values from ordinary least squares regression (OLS) for each subject to the target stimulus, we found that 20 of these features were highly sensitive to target according to a one-sample t-test ($p \leq 10^{-7}$). Among these were motor gyrus and temporal gyrus features (Figure 1), which are widely found in the literature to be sensitive to task [1, 9]. In addition, each of these features show spatial variability across time through their variable spatial maps which cannot be explained by scalar modulation of feature intensity, as linear regression tests of $\mathbf{c}_{t,n}\mathbf{p}_{t,n}^T$ to the first principle component show poor fit. Furthermore, the spatial maps that show the worst fit are often those that show sensitivity to task.

The time courses of the spatial maps, $\mathbf{c}_{t,n}\mathbf{p}_{t,n}^T$, also show sensitivity to task, including significant group differences in the β values from linear regression to the target stimuli (Figure 2). Most surprisingly, the spatial maps show significant group differentiation ($p \leq 0.05$, FDR corrected), while none of the source time courses show significant group differences. This contrasts with findings from ICA fMRI task studies [7]. This may indicate that ICA exaggerates source variation that is perhaps better accounted by IN shape, which is further supported by our finding that the spatial map variability cannot be explained by scalar modulation. However, further study is necessary to make any conclusions about spatial map variability and group differentiation.

References

[1] Elena A Allen, Erik B Erhardt, Yonghua Wei, Tom Eichele, and Vince D Calhoun. Capturing inter-subject variability with group independent component analysis of fMRI data: a simulation study. *Neuroimage*, 59(4):4141–4159, 2012.

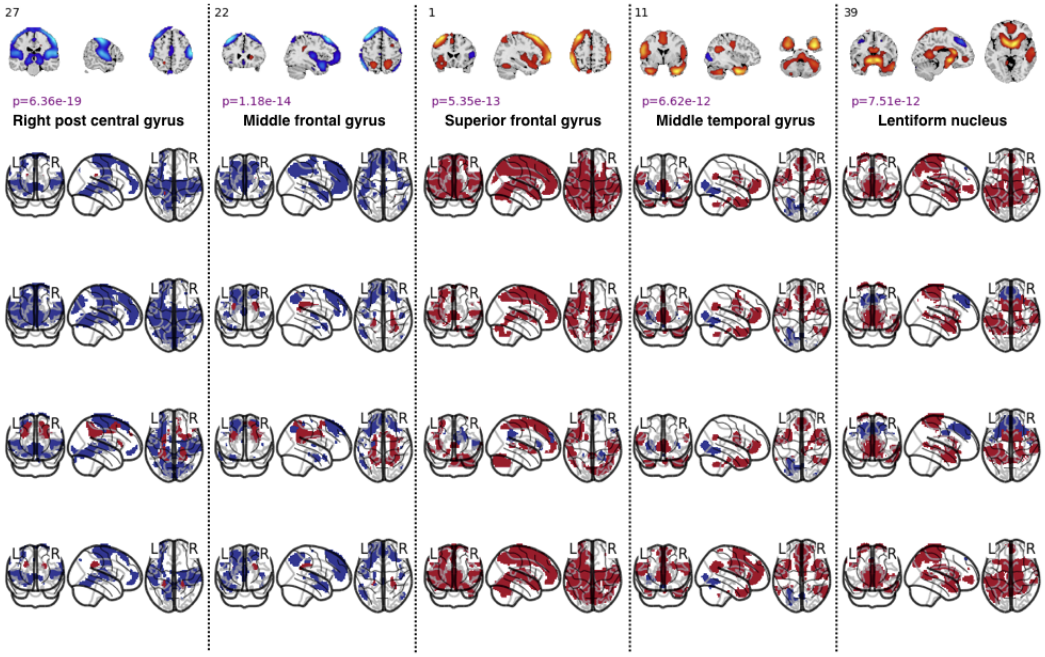

Figure 1: Five representative representative features with high sensitivity to the target stimulus. Labels are the largest region of interest. Above are the average spatial maps across time and subjects, back-reconstructed. p-values are from a one sample t-test of beta values from linear regression to the target stimulus. Below are four time steps in sequence of each feature for a single patient.

[2] Anthony J. Bell and Terrence J. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, 7(6):1129–1159, 1995.

[3] Bharat Biswal, F Zerrin Yetkin, Victor M Haughton, and James S Hyde. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magnetic Resonance in Medicine*, 34(4): 537–541, 1995.

[4] Vince D. Calhoun, Tulay Adali, Godfrey D. Pearlson, and J. J. Pekar. A method for making group inferences from functional MRI data using independent component analysis. *Human Brain Mapping*, 14, 2001.

[5] Vince D. Calhoun, Kent A. Kiehl, and Godfrey D. Pearlson. Modulation of temporally coherent brain networks estimated using ICA at rest and during cognitive tasks. *Human Brain Mapping*, 29(7):828–838, 2008.

[6] JS Damoiseaux, SARB Rombouts, F Barkhof, P Scheltens, CJ Stam, Stephen M Smith, and CF Beckmann. Consistent resting-state networks across healthy subjects. *Proceedings of the National Academy of Sciences*, 103(37):13848–13853, 2006.

[7] Wei Du, Vince D Calhoun, Hualiang Li, Sai Ma, Tom Eichele, Kent A Kiehl, Godfrey D Pearlson, and Tülay Adali. High classification accuracy for schizophrenia with rest and task fmri data. *Frontiers in human neuroscience*, 6, 2012.

[8] Geoffrey Hinton. Neural networks for machine learning. Coursera, video lectures, 2012.

[9] R Devon Hjelm, Vince D Calhoun, Ruslan Salakhutdinov, Elena A Allen, Tulay Adali, and Sergey M Plis. Restricted boltzmann machines for neuroimaging: an application in identifying intrinsic networks. *NeuroImage*, 96:245–260, 2014.

[10] R Devon Hjelm, Sergey M Plis, and Vince C Calhoun. Variational autoencoders for feature detection of magnetic resonance imaging data. *arXiv preprint arXiv:1603.06624*, 2016.

[11] Aapo Hyvärinen and Erkki Oja. Independent component analysis: algorithms and applications. *Neural networks*, 13(4):411–430, 2000.

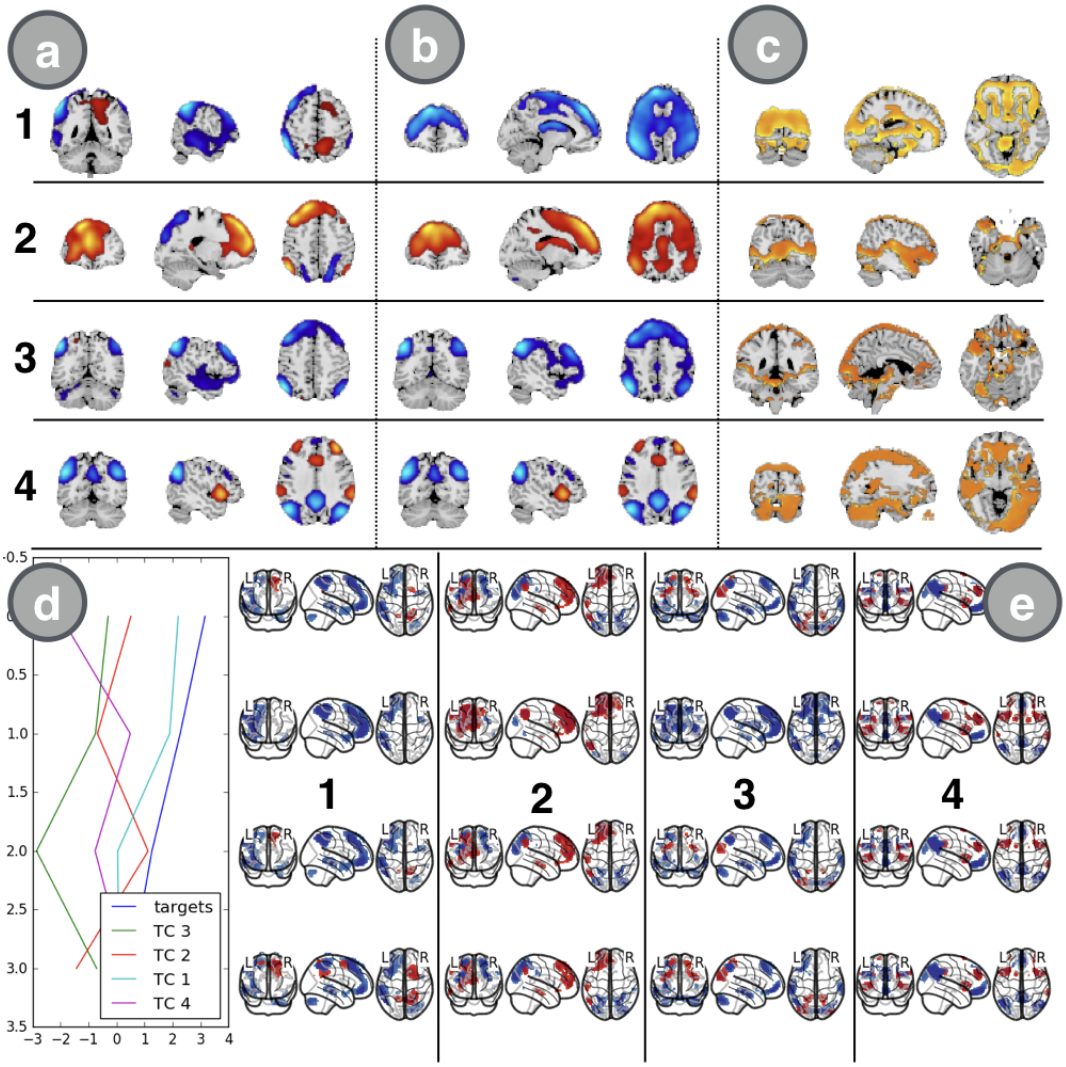

Figure 2: Four features whose spatial maps showed significant differences between healthy controls and groups. Differences were found from the betas from linear regression to the target stimulus. a) Average spatial maps back-reconstructed, b) The uniform spatial map, $\overline{\mathbf{W}}$, back-reconstructed, c) Mann-Whitney u values from significance test between healthy controls and patients, d) time courses and convolved target stimulus for e) four time steps of spatial maps for a single patient.

[12] Jack Sherman and Winifred J Morrison. Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *The Annals of Mathematical Statistics*, 21(1):124–127, 1950.

[13] Stephen M Smith, Peter T Fox, Karla L Miller, David C Glahn, P Mickle Fox, Clare E Mackay, Nicola Filippini, Kate E Watkins, Roberto Toro, Angela R Laird, et al. Correspondence of the brain's functional architecture during activation and rest. *Proceedings of the National Academy of Sciences*, 106(31):13040–13045, 2009.

[14] Nathan Swanson, Tom Eichele, Godfrey Pearlson, Kent Kiehl, Qingbao Yu, and Vince D Calhoun. Lateral differences in the default mode network in healthy controls and patients with schizophrenia. *Human Brain Mapping*, 32(4):654–664, 2011.